

Sunfish: Enabling Predictive Analytics for Datacenters Through Digital Twinning

M. Kwiatkowski Drs. D. Niewenhuis¹ Prof. A. Iosup²

¹Daily Supervisor

²Main Supervisor

Vrije Universiteit Amsterdam

June 19, 2026

Online slides: <https://www.overleaf.com/read/tknhxmqlfgtdy#87413e>

Context

21st century datacenters are primarily heterogeneous [8] and modern computational needs of AI drive managers to diversify datacenters even more [1]. In result datacenters become extremely complex and hard to operate with millions of CPU's, GPU's etc.



Figure 1.1: Society depends on datacenters to keep running, and therefore we cannot afford to let these systems break down or experience significant performance-related issues. With millions of servers in the largest datacenters, real-time management becomes very difficult. Left to right: a Google datacenter, server racks, Ada Lovelace AD102 GPU architecture.

DCDT's lack predictive analytics

We need Datacenter Digital Twins (DCDT) to be better able to detect and solve issues in critical ICT infrastructure [1]. However, DCDT's are still actively developed and lack crucial features such as predictive analytics [9] to *e.g.*, prevent unexpected failures. With predictive analysis (*e.g.*, simulation) DCDT's could save millions of lost \$USD [11].

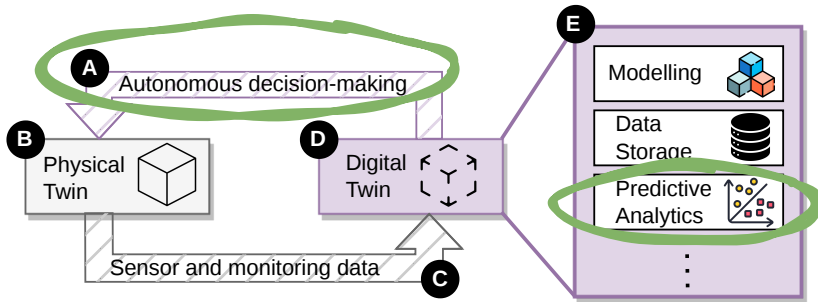


Figure 1.2: Where does our work fit within the field of datacenter digital twinning? There are 5 core elements to any Digital Twin:

A The Digital → Physical Twin link, **B** the Physical Twin (*e.g.*, the datacenter), **C** the Physical → Digital Twin link, **D** the Digital Twin, **E** the features necessary to any Digital Twin. **Highlighted areas are the contributions from this thesis, which include the autonomous actions resulting from predictive insights **A** and the predictive analysis itself within **E**.**

Main Research Question

How to enable predictive analytics for datacenters through digital twinning?

Research Question 1

How to assess the current state-of-the-art of digital twinning for datacenters?

Research Question 2

How to design a reference architecture for a predictive datacenter digital twin using discrete-event simulation?

Research Question 3

How to evaluate and validate a datacenter digital twin architecture in relation to system requirements?

Results

This is a dummy sentence meant to make the tcolorbox have more than 2 lines of text width so that I am able to show the text and the table spacing better. I hope it fits its purpose well.

Project	Simulation Technique	Focus	Stakeholders	Highlighted Features
ExaDigiT [2]	CFD/HT	Energy Loss Prediction		
SmartDC [14]	CFD/HT, BIM, AI	Heat Modelling		
DyTwin [10]	Gaussian Process Regression, ML	Anomaly Detection		
ChatTwin [7]	?	Configuration Automation		
Reducio [3]	POD	Heat and Airflow Prediction		
NetGraph [5]	Graphs	Network Management		
Kalibre [13]	ML, CFD	Heat Modelling		

Table 1.1: Comparison of selected datacenter digital twins. **Features:** 3D = 3D Visualizations; CH = Cooling/Heat, PE = Power/Energy Consumption, F = Failures, N = Network, FS = FaaS, SE = Scenario Exploration, VP = Virtual Prototyping, FD = Federation; **Tools:** AI = Artificial Intelligence, ML = Machine Learning, ODA = Operational Data Analysis; * = Predictive Analysis; ★ = Descriptive Analysis, ◆ = Prescriptive Analysis.

A generic system model

This is a dummy sentence meant to make the colorbox have more than 2 lines of text width so that I am able to show the text and the table spacing better.

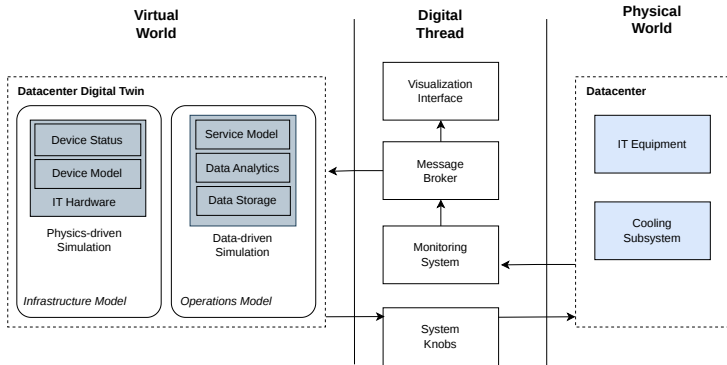


Figure 1.3: To answer RQ1 we designed a generic datacenter digital twin system model based on a comprehensive literature review and findings from Table 1.1.

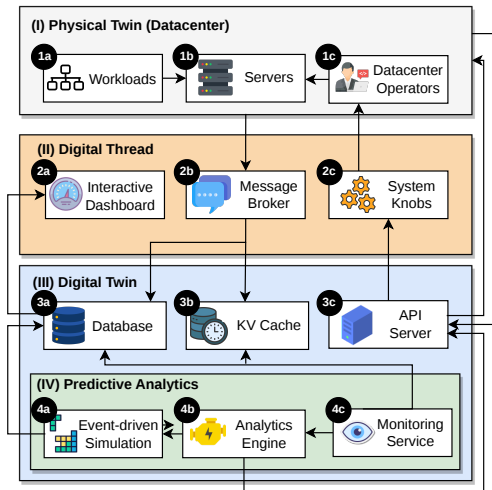


Figure 1.4: The predictive datacenter digital twin architecture.

Functional Req.

- FR1:** The system shall be able to
- FR2:** The system should be able to
- FR3:** The system needs to do this and that

Non-functional Req.

- NFR1:** The system shall be able to
- NFR2:** The system should be able to
- NFR3:** The system needs to do this and that

Main Finding I

On average, *Sunfish* achieves 12.17% less failures per task than baseline (OpenDC). Insights from predictive digital twinning yield noticeable performance difference.

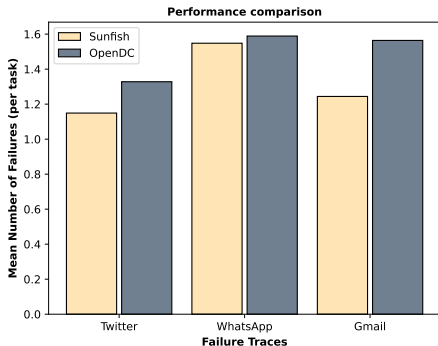


Figure 1.5: Experiment 1 – on the x-axis are different community failure traces. On the y-axis is the mean number of times a task has failed, during the entire workload. Vertical bars is standard deviation, measured over 5 repetitions.

Main Finding II

Here explain what did you find.

What is the societal context?

What problem did we solve?

How did we solve this problem?

What did we find?

What do we see in future work?

Extra Slides: References I



Jyotika Athavale, Cullen E. Bash, Wesley Brewer, Matthias Maiterth, Dejan S. Milojevic, Harry Petty, and Soumyendu Sarkar.

Digital twins for data centers.

Computer, 57(10):151–158, 2024.

URL <https://doi.org/10.1109/MC.2024.3436945>.



Wesley Brewer, Matthias Maiterth, Vineet Kumar, Rafal P. Wojda, Sedrick Bouknight, Jesse Hines, Woong Shin, Scott Greenwood, David Grant, Wesley Williams, and Feiyi Wang.

A digital twin framework for liquid-cooled supercomputers as demonstrated at exascale.

In *Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis, SC 2024, Atlanta, GA, USA, November 17-22, 2024*, page 23. IEEE, 2024.

URL <https://dl.acm.org/doi/10.1109/SC41406.2024.00029>.



Zhiwei Cao, Ruihang Wang, Xin Zhou, and Yonggang Wen.

Reducio: model reduction for data center predictive digital twins via physics-guided machine learning.

In Jorge Ortiz, editor, *Proceedings of the 9th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation, BuildSys 2022, Boston, Massachusetts, November 9-10, 2022*, pages 1–10. ACM, 2022.

URL <https://doi.org/10.1145/3563357.3564050>.



Douglas Donnellan, Andy Lawrence, and Rose Weinschenk.

Executive summary: Annual outage analysis 2025, May 2025.

URL <https://uptimeinstitute.com/resources/research-and-reports/annual-outage-analysis-2025>.



Hanshu Hong, Qin Wu, Feng Dong, Wei Song, Ronghua Sun, Tao Han, Cheng Zhou, and Hongwei Yang.

Netgraph: An intelligent operated digital twin platform for data center networks.

In *NAI'21: Proceedings of the ACM SIGCOMM 2021 Workshop on Network-Application Integration, Virtual Event, USA, August 27, 2021*, pages 26–32. ACM, 2021.

URL <https://doi.org/10.1145/3472727.3472802>.

Extra Slides: References II



Alexandru Iosup, Fernando Kuipers, Ana Lucia Varbanescu, Paola Grosso, Animesh Trivedi, Jan S. Rellermeier, Lin Wang, Alexandru Uta, and Francesco Regazzoni.

Future computer systems and networking research in the netherlands: A manifesto.

CoRR, abs/2206.03259, 2022.

URL <https://doi.org/10.48550/arXiv.2206.03259>.



Minghao Li, Ruihang Wang, Xin Zhou, Zhaomeng Zhu, Yonggang Wen, and Rui Tan.

Chattwin: Toward automated digital twin generation for data center via large language models.

In *Proceedings of the 10th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation, BuildSys 2023, Istanbul, Turkey, November 15-16, 2023*, pages 208–211. ACM, 2023.

URL <https://doi.org/10.1145/3600100.3623719>.



Dejan S. Miloijic, Paolo Faraboschi, Nicolas Dubé, and Duncan Roweth.

Future of HPC: diversifying heterogeneity.

In *Design, Automation & Test in Europe Conference & Exhibition, DATE 2021, Grenoble, France, February 1-5, 2021*, pages 276–281. IEEE, 2021.

URL <https://doi.org/10.23919/DATE51398.2021.9474063>.



National Academy of Engineering, National Academies of Sciences Engineering, and Medicine.

Foundational research gaps and future directions for digital twins.

The National Academies Press, Washington, DC, 2024.

ISBN 978-0-309-70042-9.

URL <https://nap.nationalacademies.org/catalog/26894/>

[foundational-research-gaps-and-future-directions-for-digital-twins](https://nap.nationalacademies.org/catalog/26894/foundational-research-gaps-and-future-directions-for-digital-twins).

Extra Slides: References III



Ebad Taheri, Pedro Bruel, Pavana Prakash, Gourav Rattihalli, Ninad Hogade, Aditya Dhakal, Rolando P. Hong Enriquez, Torsten Wilde, Leo Popokh, Dejan S. Milojicic, and Cullen E. Bash.

Dytwin: Federated adaptive digital twins for data centers - visualization and anomaly detection.

In *SC24-W: Workshops of the International Conference for High Performance Computing, Networking, Storage and Analysis*, Atlanta, GA, USA, November 17-22, 2024, pages 847–852. IEEE, 2024.

URL <https://doi.org/10.1109/SCW63240.2024.00120>.



Sacheendra Talluri, Leon Overweel, Laurens Versluis, Animesh Trivedi, and Alexandru Iosup.

Empirical characterization of user reports about cloud failures.

In Esam El-Araby, Vana Kalogeraki, Danilo Pianini, Fr édéric Lassabe, Barry Porter, Sona Ghahremani, Ingrid Nunes, Mohamed Bakhouya, and Sven Tomforde, editors, *IEEE International Conference on Autonomic Computing and Self-Organizing Systems, ACSOS 2021, Washington, DC, USA, September 27 - Oct. 1, 2021*, pages 158–163. IEEE, 2021.

URL <https://doi.org/10.1109/ACSOS52086.2021.00039>.



Fei Tao, Meng Zhang, Yushan Liu, and A.Y.C. Nee.

Digital twin driven prognostics and health management for complex equipment.

CIRP Annals, 67(1):169–172, 2018.

ISSN 0007-8506.

URL <https://www.sciencedirect.com/science/article/pii/S0007850618300799>.



Ruihang Wang, Xin Zhou, Linsen Dong, Yonggang Wen, Rui Tan, Li Chen, Guan Wang, and Feng Zeng.

Kalibre: Knowledge-based neural surrogate model calibration for data center digital twins.

In *BuildSys '20: The 7th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation, Virtual Event, Japan, November 18-20, 2020*, pages 200–209. ACM, 2020.

URL <https://doi.org/10.1145/3408308.3427982>.

Extra Slides: References IV



Ziting Zhang, Yu Zeng, Haoran Liu, Chaoyue Zhao, Feng Wang, and Yunqing Chen.

Smart DC: an AI and digital twin-based energy-saving solution for data centers.

In 2022 IEEE/IFIP Network Operations and Management Symposium, NOMS 2022, Budapest, Hungary, April 25-29, 2022, pages 1–6. IEEE, 2022.

URL <https://doi.org/10.1109/NOMS54207.2022.9789853>.

Extra Slides: Societal Impact

Why is this research important today?

Over 3 million jobs in the Netherlands directly depend on cloud services, which are hosted in datacenters [6]. Already the rapid expansion of datacenters has increased the presence of service failures across all cloud services [11]. We need to act now.




Figure E.1: Horizontally: the most important research areas in computer science in Netherlands. Vertically: qualities we should ensure across all research areas with the most outstanding impact on society. Datacenter manageability is a top-priority [6].

Extra Slides: Why Digital Twinning?

Definition

A DCDT mirrors the structure, context and behaviour of a datacenter [1]. The prerequisite to any digital twin is good monitoring and sensing capabilities in the physical entity. Datacenters meet this requirement easily because they already connect hundreds of monitoring sensors.



Figure E.2: Due to insufficient technological foundations, little work is available on DTs between 2003 and 2018, and it is only with the rapid growth of cloud computing, Internet-of-Things and Big Data analytics that DTs have reemerged [12]. That is why nobody used digital twins to mirror datacenters earlier.

Extra Slides: Why not pure simulation?

Predicting job failures

Preventing failure-caused outages in advance can reduce huge operational costs, as over 20% of all reported outages amount to more than 1 million US\$ [4]. Only a constant bi-directional interaction (digital twin \leftrightarrow physical entity) can achieve this.



Figure E.3: Real-time control that is tightly-coupled with the IT equipment is a prerequisite for timely predictions within seconds/minutes [1].